

MANY HAKEN HEEGAARD SPLITTINGS

ALESSANDRO SISTO

ABSTRACT. We give a simple criterion for a Heegaard splitting to yield a Haken manifold. As a consequence, we construct many Haken manifolds, in particular homology spheres, with prescribed properties, namely Heegaard genus, Heegaard distance and Casson invariant.

Along the way we give simpler and shorter proofs of the existence of splittings with specified Heegaard distance, originally proven by Ido-Jang-Kobayashi, of the existence of hyperbolic manifolds with prescribed Casson invariant, originally due to Lubotzky-Maher-Wu, and of a result about subsurface projections of disc sets (for which we even get better constants), originally due to Masur-Schleimer.

1. INTRODUCTION

Every closed connected oriented 3-manifold admits a Heegaard splitting, meaning that it can be obtained gluing two handlebodies along their boundaries. It is therefore interesting to understand what kind of information one can extract about the 3-manifold from the gluing map that defines one of its Heegaard splittings. For example, in a seminal work Hempel [Hem01] proved that if a 3-manifold is Seifert fibered or contains an incompressible torus then, for any Heegaard splitting, the distance between the disc sets of the handlebodies in the curve graph of their common boundary, which we will call Heegaard distance, is at least 2. In particular, because of geometrisation, if a 3-manifold admits a Heegaard splitting of distance at least 3 then it is hyperbolic.

In this paper, we study how the property of being Haken can be read off the data coming from a Heegaard splitting.

(Recall that the surface S embedded in the 3-manifold M admits a *compression disc* if there exists an embedding f of the 2-disc D^2 into M so that $f(\partial D^2)$ is an essential loop in S and $f(\mathring{D}^2) \subseteq M - S$. A closed connected oriented 3-manifold is *Haken* if it is irreducible and it contains an incompressible surface, i.e. an embedded connected orientable surface not homeomorphic to the 2-sphere that does not admit a compression disc.)

We give a simple criterion for a Heegaard splitting to yield a Haken manifold (Theorem 3.3). Roughly speaking, the criterion applies when there is a tight geodesic connecting the disc sets along which there are large subsurface projections at every point. We will provide an explicit construction of an embedded surface (see Definition 3.1) starting from a path in the curve

graph with certain properties (specified in Definition 2.3), and then prove that such surface does not admit compression discs.

It is relatively easy to construct splittings so that the criterion applies, and in fact there is a lot of flexibility in doing so. In particular, we can construct Haken manifolds, and in particular Haken homology spheres, that satisfy a rather long list of prescribed properties:

Theorem 1.1. *Let g, n and k be integers and suppose that either $g, n \geq 3$ or $g \geq 2, n \geq 4$. Then there exists a closed oriented 3-manifold with the following properties:*

- M is Haken,
- M is an integer homology sphere,
- M is hyperbolic,
- M has a Heegaard splitting of genus g and Heegaard distance n ,
- M has Casson invariant k .

We emphasize that various subsets of those properties were not known to be simultaneously realisable. In fact, for example, the first construction of splittings of given Heegaard distance is given in [IJK14], but these are not guaranteed to be neither Haken nor homology spheres. (There is a construction of Haken manifolds with splittings of arbitrarily large distance [Eva06], but such manifolds have positive first Betti number.) Also, the only previously known construction of hyperbolic manifolds with given Casson invariant is the one in [LMW16], where the authors do not obtain precise control on the Heegaard distance and do not show whether their manifolds are Haken or not. In fact, our construction is shorter and simpler than the ones in either of these papers, especially [LMW16], which uses probabilistic methods.

Along the way, see Corollary 2.5, we also improve the bounds and give a simpler proof of a useful result from [MS13] about subsurface projections of disc sets.

Acknowledgement. This paper would have not been possible without the contribution of Saul Schleimer, who provided precious suggestions and insights about the construction of the mapping classes in Section 4.

The author would also like to especially thank Jeff Brock for the discussions that lead to the idea for constructing incompressible surfaces, as well as Sebastian Hensel, Joseph Maher, Kasra Rafi and Juan Souto for interesting discussions.

This material is based upon work supported by the National Science Foundation under grant No. DMS-1440140 while the author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2016 semester.

2. STEADY PATHS

2.1. Background and conventions. We denote by Σ_g the closed connected oriented surface of genus g , and we will always assume $g \geq 2$.

For short, we will write “curve” instead of “essential simple closed curve” (recall that a curve is essential if it does not bound a disc or a once-punctured disc). The curve graph $\mathcal{C}(\Sigma_g)$ of Σ_g is the graph whose vertices are isotopy classes of curves on Σ_g and where two curves are connected by an edge if and only if they have disjoint representatives. With an abuse, when discussing properties of a set of curves we will implicitly assume that they are in minimal position, and when referring to the distance between two curves in $\mathcal{C}(\Sigma_g)$ we will mean the distance between their isotopy classes.

We say that a subsurface of Σ_g is non-sporadic if it is not a sphere with at most 4 discs removed or a torus with one disc removed. Similarly to the above, given an essential non-sporadic subsurface Y of Σ_g we denote by $\mathcal{AC}(Y)$ the graph whose vertices are isotopy classes of curves and essential simple arcs in Y (an arc is essential if it does not cut out a disc), with distance defined as above. For sporadic surfaces the definitions need to be adjusted; we do not recall them here since this does not play a big role in this paper, and we refer the reader to [MM99].

Curve graphs, as well as arc and curve graphs, are Gromov-hyperbolic [MM99], see also [Aou13, Bow14, CRS15, HPW15, PS15]. (This fact only plays a minor role in this paper.)

A multicurve is a collection of disjoint pairwise non-isotopic curves. Given a multicurve c , we denote $N(c)$ an open regular neighborhood.

2.1.1. Subsurface projections. We now recall some properties of subsurface projections. (The statement and proof of the criterion for being Haken do not rely on this notion, but the construction of disc sets where the criterion applies does.)

For Y an essential non-sporadic subsurface of Σ_g and a curve c in Σ_g , the subsurface projection $\pi_Y(c) \subseteq \mathcal{AC}(Y)$ is obtained as follows. First, one isotopes c so that it intersects ∂Y minimally. Then, one considers all connected components of $c \cap Y$, and defines $\pi_Y(c)$ as the set of all isotopy classes of arcs and curves that they represent. This is a set of diameter at most 1 in $\mathcal{AC}(Y)$. We will write $d_{\mathcal{AC}(Y)}(c, c')$ for $d_{\mathcal{AC}(Y)}(\pi_Y(c), \pi_Y(c'))$.

One of the fundamental facts about surface projection, and one that we will use repeatedly, is the Bounded Geodesic Image Theorem:

Theorem 2.1 (Bounded Geodesic Image Theorem, [MM00], see also [Web15]). *There exists $C \geq 0$ with the following property. Let $Y \subseteq Z$ be essential subsurfaces of Σ_g . If c_0, \dots, c_n are curves that form a geodesic in $\mathcal{AC}(Z)$ and $\pi_Y(c_i)$ is non-empty for every Y , then $d_{\mathcal{AC}(Y)}(c_0, c_n) \leq C$.*

The following is a well-known easy consequence of the Bounded Geodesic Image Theorem.

Lemma 2.2. *There exists K so that whenever c_0, \dots, c_n is a sequence of curves in an essential subsurface Y of Σ_g where consecutive curves are disjoint and not isotopic, and $d_{\mathcal{AC}(\Sigma_g - N(c_i))}(c_{i-1}, c_{i+1}) \geq K$ for all $i = 1, \dots, n-1$, then any geodesic in $\mathcal{AC}(Y)$ from c_0 to c_n contains all the c_i .*

Proof. We let $K = 2C+1$, for C as in the Bounded Geodesic Image Theorem.

We argue by contradiction. Suppose that $i \geq 2$ is minimal so that some geodesic γ from c_0 to c_i does not contain c_{i-1} (for $i = 1$ the statement is obvious, and by minimality we only have to show that γ contains c_{i-1}). Then by the Bounded Geodesic Image Theorem we have $d_{\mathcal{AC}(\Sigma_g - N(c_i))}(c_0, c_{i+1}) \leq C$. However, we can also apply the same theorem to c_0, \dots, c_{i-1} and get $d_{\mathcal{AC}(\Sigma_g - N(c_i))}(c_0, c_{i-1}) \leq C$. But then we would have $d_{\mathcal{AC}(\Sigma_g - N(c_i))}(c_{i-1}, c_{i+1}) \leq 2C$, a contradiction. \square

2.2. Definition of steady paths. We will construct surfaces in Heegaard splittings starting from paths (of multicurves) in the curve graph with certain properties described below. The key condition is a large links condition, item 5.

The notion of steady path we describe below is related to the notion of tight geodesics as defined in [MM99], and in particular for d large enough a steady path is a tight geodesic (this fact does not get used in the proof of the criterion for being Haken).

Definition 2.3. For $\mathcal{D}_0, \mathcal{D}_1$ two sets of curves on Σ_g , where $g \geq 2$, we say that a sequence of multicurves t_0, \dots, t_n is a $(\mathcal{D}_0, \mathcal{D}_1, d)$ -steady path if

- (1) the curves of the multicurve t_0 (resp. t_n) are in \mathcal{D}_0 (resp. \mathcal{D}_1),
- (2) whenever $c \in t_i, c' \in t_{i+1}$, we have that $d_{\mathcal{C}(\Sigma_g)}(c, c') = 1$,
- (3) for every $i \neq 0, 1$ (resp. $i \neq n-1, n$), every $d \in \mathcal{D}_0$ (resp. $d \in \mathcal{D}_1$) intersects t_i ,
- (4) every $d \in \mathcal{D}_0$ (resp. $d \in \mathcal{D}_1$) not in t_0 (resp. t_n) intersects $t_0 \cup t_1$ (resp. $t_{n-1} \cup t_n$),
- (5) $d_{\mathcal{AC}(\Sigma_g - N(t_i))}(t_{i-1}, t_{i+1}) \geq d$ for all $i = 1, \dots, n-1$.

2.3. Concatenations of arcs. The following lemma will be important to rule out compression discs. Essentially, the loop ℓ describes the shape of the boundary of a disc that we will encounter later in an argument by contradiction. In that context, the first condition will be guaranteed by Definition 2.3.5.

Lemma 2.4. *Let Y be a compact surface with boundary. Then there does not exist a homotopically trivial loop ℓ obtained concatenating essential arcs $\alpha_1, \dots, \alpha_k$ in Y such that:*

- *If i is even and j is odd then $d_{\mathcal{AC}(\Sigma_g - N(u))}(\alpha_i, \alpha_j) > 1$ (i.e., the arcs intersect essentially),*
- *α_i is disjoint from α_j when $i \equiv j(2)$ and $i \neq j$.*

Proof. Suppose that ℓ as in the statement exists. We can assume that the α_i are all distinct (otherwise there exists a shorter loop) and that they are in minimal position relative to their endpoints, i.e. that they do not form bigons.

Since ℓ is homotopically trivial, it can be lifted to a loop $\tilde{\ell}$ in the universal cover \tilde{Y} of Y . Denote by $\tilde{\alpha}_i$ the lifts of the α_i that concatenate to form $\tilde{\ell}$.

The key facts that we will use are that

- (1) each lift of an α_i separates \tilde{Y} ,
- (2) $\tilde{\alpha}_i$ intersects some lift of α_j in its interior if and only if $i \not\equiv j(2)$, and
- (3) if a lift of α_i intersects a lift of α_j , then it does so in one point.

(The third item follows from minimal position.)

Notice that the number of arcs k is at least 2. Let n be even so that the interior of α_n intersects $\bigcup_{m \text{ odd}} \alpha_m$ in the minimal number of points among all even n . For i odd, let $\tilde{\alpha}_n^i$ be a lift of α_n that intersects $\tilde{\alpha}_i$ in its interior. We are going to show that the $\tilde{\alpha}_n^i$ are all disjoint, which contradicts the fact that $\tilde{\ell}$ is a loop (since traveling along $\tilde{\ell}$ one crosses all the $\tilde{\alpha}_n^i$ exactly once).

Suppose that two $\tilde{\alpha}_n^i$ coincide. We can then consider distinct odd indices i, j so that $\tilde{\alpha}_n^i$ intersects $\tilde{\alpha}_j$ and $|i - j|$ is minimal. Then $|i - j| = 2$. In fact, if, say, $i \geq j + 4$ then $\tilde{\alpha}_n^{i+2}$ would intersect $\tilde{\alpha}_{j'}$ for $j' = i$ or $i + 2 < j' \leq j$, as suggested in Figure 1. More precisely, the intersection points of $\tilde{\alpha}_n^i$ with α_i and α_j lie in the same connected component of $\tilde{Y} - \tilde{\alpha}_n^{i+2}$ because distinct lifts of α_n do not intersect, while the endpoints of $\tilde{\alpha}_{i+2}$ lie in different connected components.




Figure 1. $\tilde{\alpha}_n^{i+2}$ cannot cross $\tilde{\alpha}_n^i$.

Now, it is readily seen that all lifts of α_m , m odd, that intersect $\tilde{\alpha}_{i+1}$ in its interior also intersect $\tilde{\alpha}_n^i$ in its interior. In fact, for any such lift $\tilde{\alpha}$, the intersection points of $\tilde{\alpha}_n^i$ with α_i and α_{i+2} each lie in the same connected component of $\tilde{Y} - \tilde{\alpha}$ as one of the endpoints of $\tilde{\alpha}_{i+1}$. Moreover, $\tilde{\alpha}_n^i$ also intersect in its interior two more lifts of some α_m , namely $\tilde{\alpha}_i$ and $\tilde{\alpha}_{i+2}$. From this we deduce that the interior of α_{i+1} has fewer intersections with $\bigcup_{m \text{ odd}} \alpha_m$ than the interior of α_n , a contradiction. \square

2.4. Digression: Subsurface projection of discs. We point out that Lemma 2.4, besides being a key point in the proof of Theorem 3.3 below,

also gives a significantly simpler proof of a useful lemma about subsurface projection of discs originally due to Masur-Schleimer [MS13] (which we need later). In fact, we improve the constants given by Masur-Schleimer (in this paper we measure distances between projection sets rather than diameters as in [MS13]; using this convention we should replace “1” by “3” in both conclusions below).

In the statement we use the notions of arc graph $\mathcal{A}(X)$ of a surface with boundary X , that is defined similarly to the arc and curve graph using arcs only. Distances in the arc graph can be much larger than corresponding distances in the arc and curve graph, though, so the statement in terms of the arc graph is more refined than the corresponding statement in terms of the arc and curve graph.

Corollary 2.5. (*cfr.* [MS13, Lemma 12.20]) *Let F be a compact surface with boundary (orientable or non-orientable) and let the handlebody H be the orientable $[0, 1]$ -bundle of F . Then for every essential curve d of ∂H that bounds a disc of H the following holds.*

- *If F is orientable, and hence $H = F \times [0, 1]$, let $X = F \times \{0\}$, $Y = F \times \{1\}$ and let $\tau : X \rightarrow Y$ be the involution that switches the endpoints of the fibers. Then*

$$d_{\mathcal{A}(X)}(\pi_X(d), \tau(\pi_Y(d))) \leq 1.$$

- *If F is non-orientable, let $X = \partial H - (\partial X \times (0, 1))$ and let $\tau : X \rightarrow X$ be the involution that switches the endpoints of the fibers. Then*

$$d_{\mathcal{A}(X)}(\pi_X(d), \tau(\pi_X(d))) \leq 1.$$

Moreover, $\pi_X(d)$ lies within distance 1 in $\mathcal{AC}(X)$ from a multicurve fixed by τ .

Proof. Applying an isotopy, we can make sure that d consists of a union of arcs each of which is either an essential arc of X or Y or a fiber over a boundary point of F . We can then disregard the arcs of the second type and have a sequence $\alpha'_1, \dots, \alpha'_k$ of arcs alternately in X and Y in the orientable case, and just in X in the non-orientable case. Now, the arcs $\alpha_i = \tau^i(\alpha'_i)$ concatenate to form a homotopically trivial loop (in both cases, the inclusion of X in H is π_1 -injective). By Lemma 2.4, we get that some α_i for i odd needs to be disjoint from α_j with j even. This translates into the conditions described in the statement of the corollary.

To get the conclusion about the fixed multicurve, consider some α'_i which is disjoint from some $\tau(\alpha'_j)$. Then the subsurface Z filled by α'_i and $\tau(\alpha'_i)$ is not the whole surface, and the set of essential curves in ∂Z form the required multicurve. \square

3. STEADY SURFACES

3.1. Heegaard splittings. We denote by H_g the oriented handlebody of genus g , and we identify $\Sigma_g = \partial H_g$. A *Heegaard splitting* $\mathcal{H}(\phi)$, where

$\phi : \partial H_g \rightarrow \Sigma_g$ is a homeomorphism, is the 3-manifold denote $M(\phi)$ obtained gluing two copies H_g^0, H_g^1 of H_g to the boundary components $\Sigma_g \times \{0\}, \Sigma_g \times \{1\}$ of $\Sigma_g \times [0, 1]$ using, respectively, the identity $\partial H_g^0 \rightarrow \Sigma_g$ and $\phi : \partial H_g^1 \rightarrow \Sigma_g$. We can then define the disc set \mathcal{D}_i , for $i = 0, 1$, as the set of all isotopy classes of curves on Σ_g that bound a disc in the handlebody attached to $\Sigma_g \times \{i\}$.

We say that the *Heegaard distance* of $\mathcal{H}(\phi)$ is the distance in the curve graph of Σ_g between the disc sets.

Definition 3.1. Let $\mathcal{H}(\phi)$ be a Heegaard splitting with Heegaard distance at least 2.

A *d-steady surface* in $M = M(\phi)$ is any embedded surface S in M constructed in the following way. Let t_0, \dots, t_n be a $(\mathcal{D}_0, \mathcal{D}_1, d)$ -steady path. Choose open regular neighborhoods $N(t_i)$ of the multicurves t_i , with $\overline{N}(t_i) \cap \overline{N}(t_{i+1}) = \emptyset$. Finally, let S be the union of

- a union of disjoint discs in H_g^0 (resp. H_g^1) with boundary $\partial \overline{N}(t_0)$ (resp. $\partial \overline{N}(t_n)$),
- surfaces $S_i = S'_i \times \{i/(n+1)\}$, where $S'_i = \Sigma_g \setminus (N(t_i) \cup N(t_{i-1}))$, for $i = 1, \dots, n$,
- the unions of annuli $A_i = \partial \overline{N}(t_i) \times [i/(n+1), (i+1)/(n+1)]$, for $i = 0, \dots, n-1$.

We remark that steady surfaces are orientable.

It is easy to give conditions for a steady surface to have a connected component which is not a sphere, for example:

Lemma 3.2. *Suppose that t_0, \dots, t_n is a steady path for the surface Σ_g defining the steady surface surface S . If either $g \geq 3, n \geq 3$ or $g \geq 2, n \geq 4$, and t_1, t_{n-1} consist of a single curve, then S has a connected component which is not a sphere.*

Proof. We use the notation of Definition 3.1. Consider the subsurface S' consisting of the union of all S_i with $i \neq 1, n$ and the annuli A_i for $i = 2, \dots, n-2$ (no annuli if $n = 3$). The Euler characteristic of S' is $(2-2g)(n-2)$, and S' has 4 boundary components. Under both sets of assumptions on (g, n) , we can conclude that S' has positive genus, and hence one connected component of S is not a sphere. \square

Finally, we prove the key result to construct Haken manifolds from Heegaard splittings.

Theorem 3.3. *Let $\mathcal{H}(\phi)$ be a Heegaard splitting with $g \geq 2$ of distance at least 2. If S is a 5-steady surface in $M = M(\phi)$, then ϕ does not admit a compression disc.*

Proof. We use the notation of Definition 3.1. Suppose by contradiction that the steady surface S admits a compression disc D . We now isotope D in a suitable normal form in a few steps.

First of all, applying an isotopy we can assume that the boundary of D is either

- (1) a simple loop contained in some S_i , or
- (2) a union of essential arcs in the S_i , that we call *vertical arcs*, and arcs in the A_i of the form $\{p\} \times [i/(n+1), (i+1)/(n+1)]$, that we call *horizontal arcs*.

This is just because we can remove inessential arcs in the various $\partial D \cap S_i$ and $\partial D \cap A_i$ starting from innermost ones.

Moreover, we can assume that there are finitely many simple arcs in D and simple loops in \mathring{D} (that we call intersection arcs and intersection loops) so that the intersection of \mathring{D} and $\bar{\Sigma}_g = \bigcup \Sigma_g \times \{i/(n+1)\}$ consists of the intersection loops and the interior of the intersection arcs.

Ruling out intersection loops. We now argue that we can isotope D to remove intersection loops and that case 1 does not occur. Consider an innermost intersection loop ℓ and suppose that it is contained in S_i . If ℓ bounds a disc in S_i then a simple surgery arguments allows us to replace D by a disc that has fewer loops in the intersection with $\bar{\Sigma}_g$, hence we can assume that ℓ is essential in $\Sigma_g \times \{i/(n+1)\}$ (we are not ruling out that it is parallel into the boundary of S_i , for now). Notice that the disc cut out by ℓ does not intersect one of the handlebodies and hence, when identifying $\Sigma_g \times \{i/(n+1)\}$ with Σ_g , we have $\ell \in \mathcal{D}_0 \cup \mathcal{D}_1$. Definition 2.3.3 rules out that ℓ is contained in S_i for $i \neq 1, n$, so that ℓ is contained in S_1 or S_n . But then Definition 2.3.4 implies that ℓ is parallel to a component of $t_0 \times \{1/(n+1)\}$ or $t_n \times \{n/(n+1)\}$, which bound discs in S . Hence, we can once again replace D with a disc that has fewer loops in the intersection with $\bar{\Sigma}_g$.

We can then go on and remove all loops in $\mathring{D} \cap \bar{\Sigma}_g$, and finally a very similar argument proves that case 1 does not occur because otherwise ∂D would not be essential in S .

Cutting up D . We now have that there are no intersection loops, just intersection arcs. The intersection arcs subdivide D into finitely many closed (polygonal) regions. The Euler characteristic of D equals the number of such regions minus the number of intersection arcs. We are now going to argue that each region contains at least two intersection arcs, leading to a contradiction.

In fact, suppose that a region R contains only one intersection arc α , say contained in S_i . The closure of $\partial R - \alpha$ is an arc β contained in ∂D , and hence it is a concatenation of (alternately) horizontal and vertical arcs. Moreover, the vertical arcs are all contained in either $S_i \cup S_{i+1}$ or $S_{i-1} \cup S_i$, for otherwise the interior of R would have to intersect either S_{i-1} or S_{i+1} . Since R is contained in Σ_g times an interval, we can then project ∂R to the factor Σ_g , and obtain a concatenation of paths as described in Lemma 2.4 (notice that an arc in $\Sigma_g - N(t_i)$ disjoint from t_{i-1} intersects any arc in $\Sigma_g - N(t_i)$ that intersects t_{i+1} at most once since $d_{\mathcal{AC}(\Sigma_g - N(t_i))}(t_{i-1}, t_{i+1}) \geq 5$). However,

such concatenation is homotopically trivial because we can also project R , a contradiction. \square

4. CONSTRUCTION OF HAKEN MANIFOLDS

4.1. Fixing a Heegaard splitting of S^3 .

4.1.1. *Genus at least 3.* Fix a genus $g \geq 3$. Then the handlebody H_g of genus g can be identified with the product $F \times [0, 1]$, where F is a sphere with at least 4 discs removed. We denote by σ_0, σ_1 the core curves of two connected components of $\partial F \times [0, 1]$. We claim that there exists a Heegaard splitting $H(\iota = \iota_1 \circ \iota_2)$ of the sphere S^3 so that $\iota_i \in \text{Stab}(\sigma_i)$. In fact, there is a curve c on ∂H_g that bounds a disc and separates σ_1 from σ_2 ; just consider the product of an arc in F that separates the components of ∂F corresponding to σ_1, σ_2 , and take the product with $[0, 1]$. Now, the usual gluing map that exchanges meridian and longitudes can be written as a product of two homeomorphisms that each restrict to the identity on one component of the complement of c , as required.

Denote by \mathcal{D} the disc set of H_g , so that the disc sets associated to the Heegaard splitting are \mathcal{D} and $\iota(\mathcal{D})$.

4.1.2. *Genus 2.* In genus 2 we need a slightly different construction because there is only a limited amount of “space”. The handlebody H_2 of genus 2 is homeomorphic to the orientable $[0, 1]$ -bundle of the the Möbius strip with a disc D removed, that we call F . We denote by σ_0 , the core curves of the annulus that fibers over $\partial \bar{D}$. The curve σ_0 is depicted in Figure 2 in the standard picture of the handlebody (the the orientable $[0, 1]$ -bundle of the the Möbius strip is a solid torus, and the the orientable $[0, 1]$ -bundle of F is obtained from F by “drilling a hole”). We take σ_1 to be the other curve in the picture, which admits a similar description as σ_0 in terms of another bundle.

Everything else is as above: There exists a Heegaard splitting $H(\iota = \iota_1 \circ \iota_2)$ of the sphere S^3 so that $\iota_i \in \text{Stab}(\sigma_i)$, and we denote by \mathcal{D} the disc set of H_2 .

We fix the data described in this subsection from now on.




Figure 2

4.2. Constructing large partial pseudo-Anosovs. Let \mathcal{K} be the subgroup of $MCG(\Sigma_g)$ generated by Dehn twists around separating curves. By [Mor97], there is a homomorphism $J : \mathcal{K} \rightarrow \mathbb{Z}$ so that, for all $\phi \in \mathcal{K}$, the Casson invariant of $H(\iota \circ \psi)$ is $J(\psi)$.

Lemma 4.1. *For $i = 0, 1$ the following holds. For every L there exists $\phi_i \in \mathcal{K} \cap \text{Stab}(\sigma_i)$ so that $J(\phi_i) = 1$, ϕ_i has a geodesic axis in $\mathcal{AC}(\Sigma_g - N(\sigma_i))$, and the translation distance of ϕ is at least L .*

Proof. Since J is surjective, see e.g. [LMW16, Lemma 4], we can find separating curves c_1, \dots, c_k so that the $J(\tau_{c_j})$ are coprime, where τ_{c_j} is a Dehn twist around c_j . Notice that since J is a homomorphism to an Abelian group, it takes the same value on conjugate mapping classes. In particular, without changing the set of values $\{J(\tau_{c_j})\}$, we can assume that the c_j are contained in $\Sigma_g - N(\sigma_i)$, that there are at least two of them, and that they are at least $L + 2$ apart from each other.

We now consider any product ϕ_i of sufficiently large powers of the τ_{c_j} . In order to construct a geodesic axis for ϕ_i , there is a standard procedure: One starts from geodesics connecting the c_j , takes subgeodesics connecting curves disjoint from the c_j , starts “rotating” these using the τ_{c_j} and takes concatenations. It is now easy to use the Bounded Geodesic Image Theorem, similarly to Lemma 2.2, to show that such concatenation is a geodesic line, and that the translation distance is at least L (we have not recalled the definition of the arc graph of an annulus, but the only fact about it that is needed for this construction is that the corresponding Dehn twist acts with positive translation length).

Finally, choosing the powers of the τ_{c_j} suitably, we can further ensure $J(\phi_i) = 1$, as required. \square

4.3. Moving σ_i off the disc set.

Lemma 4.2. *For $i = 0, 1$ and for every d , there exists $\phi_i \in \mathcal{K} \cap \text{Stab}(\sigma_i)$ with $J(\phi_i) = 1$ so that for every integer $k \neq 0$ we have*

$$d_{\mathcal{AC}(\Sigma_g - N(\sigma_i))}(\mathcal{D}, \phi_i^k(\sigma_{i+1})) \geq d,$$

$$d_{\mathcal{AC}(\Sigma_g - N(\sigma_i))}(\mathcal{D}, \phi_i^k \iota_i(\sigma_{i+1})) \geq d.$$

Proof. We give the proof for genus at least 3 first.

Let L, R be large enough constants to be determined later.

Let ϕ_i be as in Lemma 4.1. Let $X = F \times \{0\}, Y = F \times \{1\}$. We can conjugate ϕ_i to ensure every curve along the axis γ of ϕ_i cuts ∂X and ∂Y . By conjugating ϕ_i by a large pseudo-Anosov of X , and keeping into account that the entire axis of γ has bounded projection onto $\mathcal{AC}(X)$ and $\mathcal{AC}(Y)$, we can then ensure that, identifying X, Y with F , for every curve c on γ we have $d_{\mathcal{AC}(F)}(\pi_X(c), \pi_Y(c)) \geq R$.

Notice that $\phi_i^k(\sigma_{i+1})$ has closest point projection to γ far away from that of ∂X . If there was $c \in \mathcal{D}$ so that $d_{\mathcal{C}(\Sigma_g - N(\sigma_i))}(c, \phi_i^k(\sigma_{i+1}))$ is small,

then, by a simple Gromov-hyperbolicity argument, we would have a geodesic from $\pi_{\Sigma_g - N(\sigma_i)}(c)$ to γ that stays far from ∂X and ∂Y . Hence, $d_{\mathcal{AC}(F)}(\pi_X(c), \pi_Y(c))$ would be large, contradicting Corollary 2.5. A similar argument holds for $\phi_i^k(\iota_i(\sigma_{i+1}))$, which lies within uniformly bounded distance of $\phi_i^k(\sigma_{i+1})$.




Figure 3. Picture in $\mathcal{C}(\Sigma_g - N(\sigma_i))$.

For genus 2, the proof is similar. In this case we let X be the double cover of F contained in ∂H_2 . We can conjugate ϕ_i so that its axis has bounded projection onto $\mathcal{AC}(X)$. Furthermore, since the Möbius strip with one disc removed has bounded curve graph, we can make sure that such projection lies far away from multicurves fixed by τ , where τ is the involution described in Corollary 2.5 (in the notation of the corollary, there is a natural correspondence between multicurves fixed by τ and multicurves of F). Finally, we can show that if the projection to $\mathcal{AC}(\Sigma_g - N(\sigma_i))$ of some $d \in \mathcal{D}$ was close to either of $\phi_i^k(\sigma_{i+1})$ or $\phi_i^k(\iota_i \sigma_{i+1})$, then its projection to X would be far away from the fixed set of τ , contradicting Corollary 2.5. \square

4.4. Constructing steady paths.

Lemma 4.3. *Fix a large enough d and let ϕ_i be as in Lemma 4.2.*

Let $n \geq 2$ and let k_1, \dots, k_n be nonzero integers. Let ψ be the product, with indices modulo 2,

$$\psi = (\phi_1^{k_1} \iota_1)(\phi_2^{k_2} \iota_2) \prod_{i=3}^n \phi_i^{k_i}.$$

Then there exists a $(\mathcal{D}, \psi(\mathcal{D}), d)$ -steady path t_0, \dots, t_n with t_i consisting of a single curve for $i \neq 0, n$. Moreover, $d_{\mathcal{C}(\Sigma_g)}(\mathcal{D}, \psi(\mathcal{D})) = n + 1$.

Proof. Let $\phi'_1 = \phi_1^{k_1} \iota_1$, $\phi'_2 = \phi_2^{k_2} \iota_2$ and $\phi'_i = \phi_i^{k_i}$ for $i \geq 3$.

For $k \geq 3$, let $\psi_k = \prod_{i=1}^k \phi'_i$. Let $t_i = \psi_{i-1} \sigma_i = \psi_i \sigma_i$ for $i = 1, \dots, n$. Also, let $t_0 \subseteq \mathcal{D}$ (resp. $t_{n+1} \subseteq \psi(\mathcal{D})$) be a maximal collection of pairwise disjoint curves disjoint from t_1 (resp. t_n).

All conditions of Definition 2.3 except for item 3 can be easily verified directly. To verify item 3 we just need to observe that, provided d is large enough, for every $d \in \mathcal{D}$, every geodesic from d to t_i , $i \geq 2$, contains t_1, \dots, t_{i-1} by Lemma 2.2, and similarly for $d' \in \psi(\mathcal{D})$. Also, any geodesic from \mathcal{D} to $\psi(\mathcal{D})$ contains t_1, \dots, t_n , proving $d_{\mathcal{C}(\Sigma_g)}(\mathcal{D}, \psi(\mathcal{D})) = n + 1$ (since $d_{\mathcal{C}(\Sigma_g)}(\mathcal{D}, \psi(\mathcal{D})) \leq n + 1$ because t_0, \dots, t_{n+1} form a path). \square

4.5. Proof of Theorem 1.1. Fix g, n and k as in the statement of the theorem. By Lemma 4.3, we can construct a Heegaard splitting $H(\psi)$ so that the two disc sets $\mathcal{D}_0, \mathcal{D}_1$ in $\mathcal{C}(\Sigma_g)$ are at distance exactly n , and they are connected by a $(\mathcal{D}_0, \mathcal{D}_1, 5)$ -steady path with t_1, t_{n-1} consisting of a single curve. Hence the resulting manifold $M(\psi)$ is Haken by Theorem 3.3 and Lemma 3.2 (notice that if a possibly disconnected surface does not admit a compression disc then none of its components do). Moreover $\psi = \iota\phi$ for some $\phi \in \mathcal{K}$ (since \mathcal{K} is normal), so that $M(\psi)$ is an integer homology sphere and by choosing the exponents k_i in Lemma 4.3, we can make sure that $J(\phi) = k$, i.e. that the Casson invariant of $M(\psi)$ is k .

Since the Heegaard splitting has distance at least 3, the resulting manifold is hyperbolic by results in [Hem01] and Thurston's hyperbolisation. \square

REFERENCES

- [Aou13] Tarik Aougab. Uniform hyperbolicity of the graphs of curves. *Geometry & Topology*, 17(5):2855–2875, 2013.
- [Bow14] Brian H. Bowditch. Uniform hyperbolicity of the curve graphs. *Pacific J. Math*, 269(2):269–280, 2014.
- [CRS15] Matt Clay, Kasra Rafi, and Saul Schleimer. Uniform hyperbolicity of the curve graph via surgery sequences. *Algebraic and Geometric Topology*, 14(6):3325–3344, 2015.
- [Eva06] Tatiana Evans. High distance Heegaard splittings of 3-manifolds. *Topology Appl.*, 153(14):2631–2647, 2006.
- [Hem01] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001.
- [HPW15] Sebastian Hensel, Piotr Przytycki, and Richard C. H. Webb. 1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs. *J. Eur. Math. Soc. (JEMS)*, 17(4):755–762, 2015.
- [IJK14] Ayako Ido, Yeonhee Jang, and Tsuyoshi Kobayashi. Heegaard splittings of distance exactly n . *Algebr. Geom. Topol.*, 14(3):1395–1411, 2014.
- [LMW16] Alexander Lubotzky, Joseph Maher, and Conan Wu. Random methods in 3-manifold theory. *Proceedings of the Steklov Institute of Mathematics*, 292(1):118–142, 2016.
- [MM99] Howard A Masur and Yair N Minsky. Geometry of the complex of curves I: Hyperbolicity. *Inventiones mathematicae*, 138(1):103–149, 1999.
- [MM00] Howard A Masur and Yair N Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geometric and Functional Analysis*, 10(4):902–974, 2000.
- [Mor97] Shigeyuki Morita. Casson invariant, signature defect of framed manifolds and the secondary characteristic classes of surface bundles. *J. Differential Geom.*, 47(3):560–599, 1997.
- [MS13] Howard Masur and Saul Schleimer. The geometry of the disk complex. *J. Amer. Math. Soc.*, 26(1):1–62, 2013.
- [PS15] Piotr Przytycki and Alessandro Sisto. A note on acylindrical hyperbolicity of mapping class groups. *Proceedings of the 7th MSJ-SI, Hyperbolic Geometry and Geometric Group Theory*, to appear, 2015.
- [Web15] Richard C. H. Webb. Uniform bounds for bounded geodesic image theorems. *J. Reine Angew. Math.*, 709:219–228, 2015.